

Description

System and Method for Analyzing the Performance of Multiple Transportation Streams of Streaming Media in Packet-Based Networks

BACKGROUND OF INVENTION

FIELD OF INVENTION

[0001] The present invention relates generally to the field of streaming. More specifically, the present invention is related to analyzing streaming data in packetized form.

DISCUSSION OF PRIOR ART

[0002] Many electronic networks such as local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs) are increasingly being used to transport streaming media whose real-time data transport requirements exhibit high sensitivity to data loss and delivery time distortion. The technical literature is replete with various schemes to implement Quality of Service (QOS) on

such networks to address the requirements of streaming media, especially when intermixed with conventional, time-insensitive, guaranteed delivery protocol stack data traffic. Furthermore, for efficiency reasons, the streaming media transport often uses a non-guaranteed delivery upper layer protocol stack such as UDP/IP making recovery of data in the presence of packet loss difficult. Regardless of whether QOS-enabled or non-QOS-enabled networks are employed, it is necessary to monitor the behavior of packet loss, delivery time distortion, and other real-time parameters of the network to assure satisfactory quality streaming media delivery.

[0003] There exists a variety of defined Management Information Bases (MIBs) which include definitions for a number of network parameters such as packet loss, inter-arrival times, errors, percentage of network utilization, etc., whose purpose is to indicate to a network manager the general operating conditions of the network. Such traditional forms of monitoring network behavior cannot easily indicate the effects that network performance has on a single or a group of individual streaming media streams. Data gathering from MIBs operating across a range of network layers combined with a highly skilled and experi-

enced practitioner would be required to simply determine the jitter imposed on a single MPEG video stream, for instance, and would only be possible by post-processing data gathered while the network was in operation. Determining the cause of a fault in a streaming media stream may be possible through such analysis but lacks the real-time indication of a network fault that is required to maintain high-quality networks such as for video or audio delivery. It also does not address the need to monitor large numbers of streams in real-time such as streams of Video-on-Demand (VoD) networks using less technically skilled operations personnel, as would be necessary to enable implementation of continuous cost-effective quality control procedures for widely deployed networks such as for VoD.

[0004] Histograms are often used in prior art schemes to present the arrival time behavior of packets on a network, but such histograms only represent the aggregate behavior of packets arriving at the measurement node due to the need to combine MIB data from a range of network layers to extract sufficient information to track a particular stream's performance. Traditional histograms define the jitter between any two packets. Streaming media requires more

in-depth knowledge, such as the time variation across many packets referred to as the "network jitter growth". This network jitter growth affects the streaming media quality as experienced by the user due to intermediate buffer overflow/underflow between the media source and its destination.

[0005] Network jitter growth of a media stream due to traffic congestion can also be an indicator of an impending fault condition and can thus be used to avoid transport failures rather than simply to react to faults after they occur. Conventional post-processed MIB analysis is inadequate for these purposes as described above.

[0006] The concept of regulating stream flow in a network based on the leaky bucket paradigm describes a methodology that might be used to prevent intermediate buffer overflow and packet jitter by regulating the outflow of data based on a set of parameters configured to optimize a particular flow. This does not address the need to analyze and continuously monitor multiple streams as is required during the installation and operation of networks carrying streaming media, especially for those enterprises whose revenue is derived from the high quality delivery of streaming media, such as broadcast and cable television

entities.

[0007] A common prior art scheme used to effectively monitor multiple video streams is to decode each stream's MPEG content (for the video example) and display the streams on a large group of television screens. Monitoring personnel then watch the screens looking for any anomalous indications and take appropriate corrective action. This is a highly subjective and error prone process, as there is a possibility that a transient fault might be missed. This is also a reactive process, as corrective action can only be taken after a fault has occurred. Furthermore, this is also an expensive process in terms of both equipment and personnel costs. It also provides little or no indications of the root cause of the fault, thus adding to the time required for implementing corrective action. This approach also does not easily scale to modern video delivery systems based upon emerging, cost-effective high-bandwidth, networks intended to transport thousands of independent video streams simultaneously. In addition, this approach cannot pinpoint the location of the fault. To do so, the personnel and equipment must be replicated at multiple points in the distribution network, greatly increasing the cost. For this to be effective, the personnel

must monitor the same stream at exactly the same time for comparison.

[0008] Many types of network delivery impairments are transient in nature affecting a limited number of packets during a period of momentary traffic congestion, for example. Such impairments or impairment patterns can be missed using traditional monitoring personnel watching video monitors. By not recognizing possible repeating impairment patterns, faults can exist for much longer periods because after the fault has passed, there is no residual trace information available for analysis. The longer a fault persists, the worse the customer satisfaction levels, and the greater the potential for lost revenues.

[0009] Whatever the precise merits, features, and advantages of the above-mentioned prior art schemes, they fail to achieve or fulfill the purposes of the present invention.

SUMMARY OF INVENTION

[0010] The present invention provides for a system and method for analyzing packetized network traffic. In one embodiment, the system comprises: (a) one or more interfaces to forward a copy of the network traffic comprising one or more streams; (b) one or more filters to receive and filter the forwarded network traffic to isolate at least one

stream; and (c) a native streaming interface to receive packetized data corresponding to the isolated stream(s), wherein the native streaming interface provides minimum time distortion to permit media stream analysis and monitoring to indicate the network's influence on the isolated stream(s) and measure each isolated stream's conformance to a pre-determined stream standard.

[0011] In one embodiment, the system for analyzing packetized network traffic comprises: (a) a compute engine to compute statistics associated with an isolated stream, wherein the statistics for each stream comprise at least a delay factor (DF) defining an instantaneous flow rate balance representing a virtual buffer delay that is needed to prevent data loss and absorb network jitter growth; and (b) one or more interfaces to forward the computed statistics for each streams of interest to a data consumer.

[0012] In another embodiment, the present invention provides for a system and method for analyzing packetized network traffic comprising one or more transportation streams. The system comprises: (a) one or more network interfaces to receive streaming network traffic associated with the transportation streams; (b) one or more filters to filter one or more streams of interest in the received

transportation streams; (c) a compute engine comprising one or more finite state machines to compute index values associated with the streams of interest, wherein the index values for each stream comprising at least: a delay factor (DF) and a media loss rate (MLR); and (d) one or more interfaces to forward the computed index values for the streams of interest to a data consumer.

[0013] In another embodiment, the present invention's method comprises the steps of: (a) receiving network traffic comprising one or more transportation streams; (b) filtering the received traffic and isolating a transportation stream from the transportation streams; (c) computing statistics associated with the isolated transportation stream, wherein the statistics comprise at least a delay factor (DF) and a media loss rate (MLR); and (d) forwarding the computed statistics to a data consumer.

[0014] The DF value defines an instantaneous flow rate balance representing a virtual buffer delay that is needed to prevent data loss and absorb network jitter growth, and the MLR value represents the number of media packets lost or corrupted.

BRIEF DESCRIPTION OF DRAWINGS

[0015] Figures 1a–c illustrate several methods of tapping an ex–

isting network traffic flow via the present invention's computing element.

[0016] Figure 2 illustrates one embodiment of the present invention's computing element which analyzes network traffic.

[0017] Figure 3 illustrates an extended embodiment of the present invention wherein a controller is used for controlling the computing element.

[0018] Figure 4 illustrates another extended embodiment of the present invention wherein an encoder is used to encode the statistics calculated by the computing engine.

[0019] Figure 5 illustrates an internal block diagram of the computing element and its interconnection with the control and logging system.

[0020] Figure 6 illustrates an adder and a counter that form a part of the compute engine.

[0021] Figure 7 illustrates a method associated with the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0022] While this invention is illustrated and described in a preferred embodiment, the invention may be produced in many different configurations. There is depicted in the drawings, and will herein be described in detail, a preferred embodiment of the invention, with the understand-

ing that the present disclosure is to be considered as an exemplification of the principles of the invention and the associated functional specifications for its construction and is not intended to limit the invention to the embodiment illustrated. Those skilled in the art will envision many other possible variations within the scope of the present invention.

[0023] Many streaming media systems, such as VoD, broadcast television control centers, or satellite-based video distribution operations utilize packetized data networks for their low-cost and omnipresence in modern data systems. The present invention monitors these existing network conduits by sampling the data contained therein with minimal alteration of its characteristics.

[0024] Figures 1a-c illustrate several methods of tapping an existing network traffic flow via the present invention's computing element 105. Figure 1a illustrates a setup wherein an ordinary network switch or router 103, which, while performing packet switching or routing on the traffic from its many ports, such as 101 and 102, also provides for a "mirror" or "monitor" port 104. Port 104 makes all data from a desired port available to the present invention's computing element 105. Alternatively, as shown in Figure

1b, a passive network tap 108 diverts a portion of the network traffic flow energy from one network port 106 to the other network port 107 and transmits that portion via a port 109 to the present invention's computing element 105. Figure 1c illustrates yet another method to tap the existing network flow via inserting the present invention's computing element 105 directly in-line with the network link to be observed via network ports 110 and 111.

[0025] In the examples of Figures 1a–b, the computing elements 105 used in each case are identical. In the example of Figure 1c, the computing element 105 also actively forwards all traffic from network connection 110 to network connection 111 and vice versa, while simultaneously providing all traffic to the equivalent internal functionality of the computing elements designated 105.

[0026] Figure 2 illustrates one embodiment of the present invention's computing element 200 which analyzes network traffic 202. Computing element 200 comprises at least one network interface 204 to receive network traffic, one or more filters 206 to filter the received network traffic, at least one computing engine 208 to compute network statistics associated with the filtered network traffic via one or more finite state machines 210, and at least one

network interface 212 to accept control instructions and transmit the computed statistics to a data consumer. Network interface 204 interfaces with the network link to be monitored via network connections 203. Network link protocols that support such packet-based transmission include, but are not limited to, 802.3 (Ethernet), 802.4, 802.5, USB, ATM, SONET, 802.11, Fibre-channel, Firewire or 1394, Infiniband, Bluetooth, 802.11, 802.15, 802.16, 802.17, ZigBee, or a native streaming video interface such as DVB-ASI.

[0027] The streaming media traffic of interest, which may consist of many individual streams of traffic, is filtered (via one or more filters 206) from the incoming network traffic 202 and processed by the finite state machines 210 of computing engine 208 to reduce its measured transmission characteristics to a set of statistics or critical parameters known as an "Index". The Index can be communicated to a logging system with alarm values set for convenient human monitoring. For example, warnings can be forwarded to a data consumer when the computed statistics exceeds a predetermined threshold or rate-of-change. It should be noted that one computing engine can be used to track several streams of interest. Similarly, one or more com-

puting engines can be used to track several streams of interest. Hence, the number of computing engines or the number of streams to be tracked should not be used to limit the scope of the present invention.

[0028] In one preferred embodiment, the Index, known as the Media Delivery Index (MDI) consists of two parts: the Delay Factor (DF) and the Media Loss Rate (MLR). This embodiment is especially valuable for constant bit rate MPEG-2 Transport Streams carried over a network such as a packetized network. The DF represents the Instantaneous Flow Rate Balance (IFRB) and is derived in the computing element. The MLR represents the number of lost or corrupted media packets and is readily derived from tracking the Continuity Counter (CC) for the MPEG-2 transport stream application or from a sequence counter or the like for protocols, such as RTP, which support the same. The MDI (DF:MLR) then represents the two key factors which describe the dynamic behavior of streaming media over packetized networks: packet jitter growth and packet loss. This Index provides at-a-glance determination of traffic impairment as well as an indication of the operating margin of a network. By modifying the calculation of the IFRB, the DF may also be used with variable bit

rate streaming media transport over packetized networks.

[0029] Figure 3 illustrates an extended embodiment of the present invention wherein a controller 302 is used for controlling the computing element 200. Controller 302 transmits, via an interface, control instructions from a management system to modify system-level state-based logic data associated with the computing element 200, and receives, via the interface, the analysis results generated by the computing element 200.

[0030] Figure 4 illustrates another extended embodiment of the present invention wherein encoder 402 is used to encode the statistics calculated by computing engine 208. Then, the encoded statistics 404 is transmitted to a data consumer via one or more interfaces 212. Some examples of encoding include (but are not limited to) encryption (such as for security), compression, or code format conversion (e.g., convert data in an ASCII format for readability).

[0031] It should be noted that more than one network interface can be used to receive network traffic. For example, Figure 5 illustrates computing element 105 (as used in Figure 1c) with two network interfaces 516 and 517, wherein computing element 105 is used for analyzing one or more streaming media flows. The two network interfaces 516

and 517 interface with the network link to be monitored via network connections 514 and 515. As in Figure 2, network link protocols that support such packet-based transmission include, but are not limited to, 802.3 (Ethernet), 802.4, 802.5, USB, ATM, SONET, 802.11, Fibre-channel, Firewire or 1394, Infiniband, Bluetooth, 802.11, 802.15, 802.16, 802.17, ZigBee, or DVB-ASI. In operation, data received from network connection 515 is decoded via network interface 517 and the resulting data is forwarded to the filter and compute engine 520 and to the other network interface 516. Then, network interface 516 forwards the data to the network connection 514, thus completing the connection from network interface 515. Thus, all data received from network interface 515 is forwarded to network interface 514 with a minimum of distortion while making all the same data available for analysis by other components of the computing element. Likewise, all data from network connection 514 is forwarded to network connection 515 while also being forwarded to the filter and compute engine 520. The result is a continuous full duplex connection between network connections 514 and 515 providing an uninterrupted network traffic flow while simultaneously providing all network data to the filter and

compute engine 520. Alternatively, as per figure 1a and figure 1b, the computing element 105 may require only a single network interface, but otherwise performs as described above, with network data being forwarded to the filter and compute engine 520.

[0032] The filter and compute engine 520 is configured via interface 521 such that it can filter the desired streaming media flows from other network traffic types for further analysis. For example, to analyze MPEG-2 streaming video over UDP/IP protocols, the filter can be configured to accept only layer-2 packets with the IP protocol type and only IP frames with UDP protocol types and only UDP datagrams that encapsulate MPEG-2 transport streams. After performing the appropriate filtering function, the compute engine calculates the components that comprise the Index value for a given streaming media flow. The Index values, and other statistics regarding the flow, are forwarded to the network interface 522 via interface 521. Then, interface 523 is used to convey the Index values to a data consumer such as an application running, for example, in a workstation consisting of control software and a logging system 524, collectively referred to as a "management" system. Network Interface 522 need not be the same type as 516 or

517 (i.e., a RS-232 serial port). Its bandwidth via the choice of physical and link layer protocols may be scaled or sized to match the amount of data expected to be handled. It should be noted that network interface 522, interface 523, and workstation (management system) 524 may be physically co-located with the computing element 105 and need not be external.

[0033] In one embodiment, the compute engine comprises at least one finite state machine counter as shown in Figure 6. The finite state machine counter is used to compute an Instantaneous Flow Rate Balance (IFRB). Counter 628 is loaded when a packet has been received via 627. The counter is loaded with the sum of the current count and the number of bits received in this packet 625 from the adder 626. Counter 628 decrements its count at each clock input pulse 629 whose rate is set to the nominal streaming media rate. Further, counter 628 is cleared at any time via the 632 clear signal. The counter output 630 indicates the number of bits that have been received at the point of test but not yet consumed, assuming that a virtual terminal device which consumes or "uses" the streaming media flow (such as a video decoder for a streaming video media case) drains the data received at a nominal media rate at

this network location. Thus, the counter output 630 represents the size of a buffer that would be needed to prevent data loss and absorb the network jitter growth due to data arriving via a packetized network. It should be noted that counter 628 may also result in negative numbers during periods between a burst of data thus representing the size of a virtual terminal's buffer needed to be prefilled to avoid underflow. Adder 626 and counter 628 may also be combined into a single entity to simply track the net difference between bits received on the packetized network side and the bits out based upon an expected drain rate. The actual quantity being tracked may be bits or any derivative thereof (bytes, words, etc.). It is important to note that the bits counted are only those subject to the drain rate. Typically, this is the payload of the packet (i.e., no headers or overhead.) For example, in the case of an MPEG-2 transport stream sent via Ethernet IP/UDP, the bits tracked would typically be the MPEG-2 transport stream packets contained within the Ethernet frame, excluding the IP/UDP headers and Ethernet CRC. The present invention further extends to using streaming media streams that are variable bit rate in nature. Variations in media bit rate may be accommodated by monitoring and

updating the expected drain rate used in IFRB calculation along with the stream. Since this finite state machine is simple, it can operate at common media rate speeds and can be replicated easily and compactly if implemented in hardware such as an FPGA, ASIC, or discrete logic, making possible an array of such machines such that one may be dedicated to each streaming media flow. Furthermore, the filter and compute engine can also be configured to capture and track other streaming media flow parameters of interest such as an MPEG-2 transport stream's continuity counters to detect dropped or corrupted packets, stream identifiers, etc.

[0034] Returning to the discussion of Figure 5, streaming media flow parameters as described above can be forwarded via a network Interface 521, and network connection 522, and external network 523, or via any type data interface as they are captured or buffered in a memory in the filter and compute engine for later retrieval by a workstation 524. In some instances, the streaming media content itself may be presented to the workstation 524 via the same path for additional analysis. They may be combined with a time stamp at either the filter and compute engine 520 or the workstation 524. Long term logs may be maintained by 524

for trend analysis, coincident analysis with other network events, the start and end of particular streaming media flows, etc. Alternatively, workstation 524 can show an instantaneous view of streaming media parameters for human monitoring. High and low watermark values may be set in the computing element 105 or in the workstation 524 for the Index parameter or any measured parameter, such that if exceeded, will be logged or trigger an alarm; this functionality may be used to warn of possible impending faults such as deviations from nominal in the flow rates that could cause a network or terminal device buffer to overflow or underflow. The Index value indicates the network's instantaneous operating jitter margin. Additionally, the rate of sampling of such parameters can be reduced to decrease the load on interface 523 during benign network conditions or increased to provide a more detailed analysis of an identified fault. Either the computing element or workstation 524 may produce long term analysis as well by performing additional computational operation on the IFRB.

[0035] In some instances, workstation 524 functionality may be integrated with the filter and compute engine for a direct display of information to the user.

[0036] It should be noted that a pure hardware, a pure software, and a hybrid hardware/software implementation of the filter and compute engine components is envisioned and should not be used to limit the scope of the present invention.

[0037] It should be noted that various kinds of interfaces can be used for establishing a packet-based communication session between the external interfaces (514 or 515 or 523) and the computing element, such as (but not limited to) a gigabit Ethernet network controller or a 10/100 Mbit/s Ethernet network interface card. Moreover, one skilled in the art can envision using various current and future interfaces and, hence, the type of packetized network interface used should not be used to limit the scope for the present invention.

[0038] In one embodiment, bandwidth for the transportation of network parameters via interface 523 as discussed above is allocated in an "on-demand" fashion, wherein full channel (network conduit) bandwidth is allocated and available to the data consumer. Compute engine 520 can track nearly any set of parameters or events, such as the last N-packets received or statistics acquired, storing it in a circular buffer. Thus, when a critical event occurs such as

streaming media data loss, bandwidth would be allocated "on-demand" to report the tracking information leading up to the critical event to the workstation analysis device 524 through the interface 523. Having pertinent information about what traffic the network was handling (not only at the time of the critical event but leading up to it as well) presented "on-demand" at the time of the critical event is very powerful. Having this information greatly reduces the "hunting" time required to identify the cause of the critical event. This information could be gathered remotely as well, given a suitable network type for 523. Expanding on the "on-demand" possibilities for parameter reporting, bandwidth may also be allocated "on-demand" on either network interfaces 514 or 515 in an in-band reporting fashion, facilitating the monitoring by equipment on the same distribution network as the streaming media.

[0039] If the network Interface 523 is an ASI (Asynchronous Serial Interface, as in DVB-ASI) type and the streaming media content itself is presented to the Interface in such a way as to minimize instrument timing distortions, a conventional streaming media specific analyzer or monitor may be utilized to not only measure the stream's conformance to expected stream standards but also to indicate the in-

fluence of network behavior. In this configuration, the computing element may be thought of as a protocol converter as well.

[0040] The present invention's system can be used in debugging various embedded systems within the streaming media's transport network. Various equipment utilized in the transportation or creation of the streaming media may allow debugging and/or parameter manipulation via the transport network as well as provide its own statistical operational information (i.e., its own system "health"). This makes possible the cross-correlation of the system's overall state/health. The invention acquires such control information via a network channel and may use its filter and compute engine capabilities to provide either the raw or processed data to a Workstation Monitor/Logger as described for Index data above.

[0041] The present invention allows the implementer the ability to scale the amount of in-band or out-of-band measured or sampled data to pass through the system up to the maximum supported by the network conduit and down to nothing. Additionally, the present invention provides the ability to scale with improvements in network conduit technology. For example, the faster the network conduit,

the more measurements or sampled data can pass. Moreover, as high-speed systems continue to evolve, their network conduit's bandwidth is usually increased proportionately to facilitate the use of the high-speed system itself (i.e., a faster network conduit is part of the main feature-set of the system; bandwidth is thereby increased by necessity). The present invention accommodates such increases in bandwidth associated with the network conduit and utilizes such high-speed systems to extract measurements or sampled data at a faster rate.

[0042] Figure 7 illustrates a method 700 associated with an embodiment of the present invention. In step 702, network traffic is received by a network interface, wherein the traffic comprises one or more streams of packetized data. Next, in step 704, the received traffic is filtered to isolate at least one stream of packetized data. In step 706, an Index is computed for the filtered stream of packetized data. In one preferred embodiment, the Index, known as the Media Delivery Index (MDI), consists of two parts: the Delay Factor (DF) and the Media Loss Rate (MLR). The DF represents the Instantaneous Flow Rate Balance (IFRB) and is derived in the computing element as described earlier. The MLR represents the number of lost or corrupted me-

dia packets and is readily derived from tracking the Continuity Counter (CC) for the MPEG-2 transport stream application or from a sequence counter or the like for protocols, such as RTP, which support the same. The MDI (DF:MLR) then represents the two key factors which describe the dynamic behavior of streaming media over packetized networks: packet jitter growth and packet loss. This Index provides at-a-glance determination of traffic impairment as well as an indication of the operating margin of a network. Then, in step 708, the computed statistics are forwarded to a data consumer, such as one running in a workstation. In one embodiment, a quality of service (QOS) metering scheme is implemented based upon adjusting traffic priority between the forwarded computed network statistics and the streaming network traffic.

[0043] Furthermore, the present invention includes a computer program code-based product, which is a storage medium having program code stored therein which can be used to instruct a computer to perform any of the methods associated with the present invention. The computer storage medium includes any of, but not limited to, the following: CD-ROM, DVD, magnetic tape, optical disc, hard drive,

floppy disk, ferroelectric memory, flash memory, ferromagnetic memory, optical storage, charge coupled devices, magnetic or optical cards, smart cards, EEPROM, EPROM, RAM, ROM, DRAM, SRAM, SDRAM, and/or any other appropriate static or dynamic memory or data storage devices.

[0044] Implemented in computer program code-based products are: (a) receiving network traffic comprising one or more transportation streams; (b) filtering the received traffic and isolating a transportation stream from the transportation streams; (c) computing statistics associated with the isolated transportation stream comprising at least a delay factor (DF) and a media loss rate (MLR), wherein DF defines an instantaneous flow rate balance representing a buffer size that is needed to prevent data loss and absorb network jitter growth, and MLR represents the number of media packets lost or corrupted; and (d) forwarding the computed statistics to a data consumer.

CONCLUSION

[0045] A system and method has been shown in the above embodiments for the effective implementation of a system and method for measuring and exposing the dynamic behavior of streaming media over a packet-based network.

While various preferred embodiments have been shown and described, it will be understood that there is no intent to limit the invention by such disclosure but, rather, it is intended to cover all modifications and alternate constructions falling within the spirit and scope of the invention as defined in the appended claims. For example, the present invention should not be limited by the number of network interfaces, number of filters, number of streams handled by the compute engine, type of packetized network conduit, location of control software, choice of hardware or software implementation of bandwidth provisioning or filter or compute engine, type of streaming media data, choice of hardware or software implementation of the "on-demand" embodiment, computing environment, or specific hardware associated with the network interfaces, filter device, or compute engine system.

[0046] The above systems are implemented in various computing environments. For example, the present invention may be implemented on a conventional IBM PC or equivalent, multi-nodal system (e.g., LAN) or networking system (e.g., Internet, WWW, wireless web). All programming and data related thereto are stored in computer memory, static or dynamic or non-volatile, and may be retrieved by the user

in any of: conventional computer storage, display (i.e., CRT) and/or hardcopy (i.e., printed) formats. The programming of the present invention may be implemented by one skilled in the art of computer systems and/or software design.